Table of Contents

Supercomputing '94 Special Issue
Volume 2, Number 7

- NAS Tutorials and Papers at Supercomputing '94
- NAS CFD and WWW Demos at Show
- Stereoscopic Solutions to CFD Problems
- p2d2: Portable Parallel/Distributed Debugger
- Updates on NAS Projects
- JOVE Researcher Develops Algorithms for SP2
- Virtual Eye Surgery
- Credits for this Issue
- This issues front cover

Next Article | Contents | Main Menu | NAS Home

NAS Offers Tutorials and Papers at Supercomputing '94

by Marcia Redmond

An estimated five thousand people will converge on Supercomputing '94 in Washington, D.C., November 14-18. This year's conference focuses on research and education in computational science and engineering, in addition to serving its traditional purpose of advancing the science and application of supercomputing technology.

Members and affiliates of the NAS Systems Division will again present several tutorial sessions and papers. Titles, authors, and brief descriptions are listed below. Refer to the Supercomputing '94 program for days and times.

Tutorials:

The Science of Benchmarking

David Bailey (NAS) and Roger Hockney (Reading University, PA) present a scientific approach to benchmarking.The authors present a clear set of units and symbols, followed by a carefully defined set of performance parameters and metrics, and finally a hierarchy of parallel benchmarks by which to measure them.

PVM/HeNCE: Tools for Heterogeneous Network Computing

Vaidy Sunderam (NAS Research Grant, Emory University) presents an overview, which focuses on developing concurrent applications for PVM, using several models of parallelism. He also discusses debugging and profiling methods, common programming and operational errors, and performance analysis and tuning. He describes the use of HeNCE in graphically assembling concurrent applications from simple building blocks, and demonstrates the applications' profiling and execution visualization capabilities.

High Performance Computing for Scientific Applications: An Introduction

Subhash Saini (right) and Horst Simon, both of NAS, offer guidance in the rapidly growing field of scientific computing using massively parallel supercomputers. Saini and Simon also present their actual experiences with the Thinking Machines CM-5 and the Intel Paragon at NAS, along with experiences from NASA re searchers on the CRAY T3D and the IBM SP2.

ATM in a Supercomputer Network Environment

Jim McCabe (NAS) presents an overview of Asynchronous Transfer Mode network technology. He discusses factors leading to the development of ATM, existing and potential applications for ATM-based supercomputer networks, and the evolution towards services-based networking.

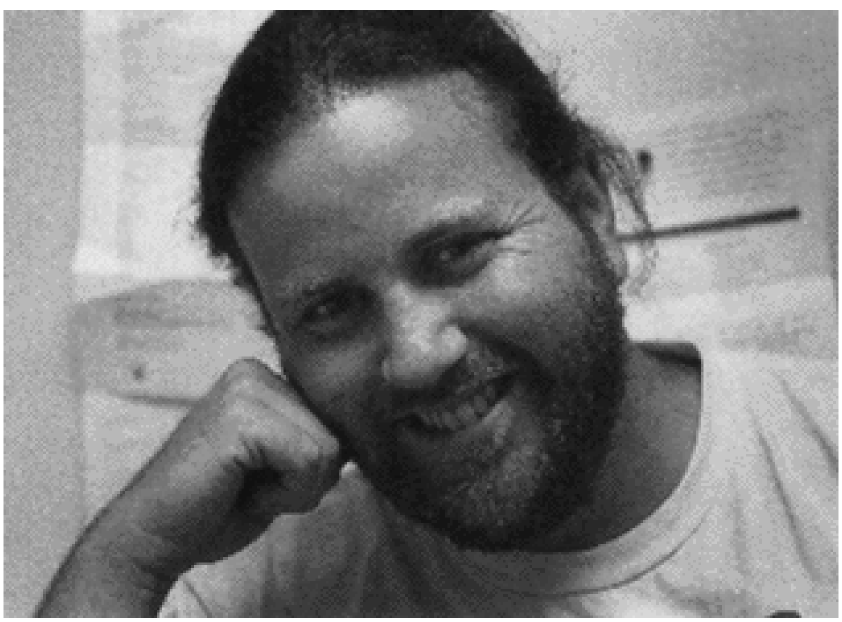

Sorting Out Communication Libraries: A Comparison of NX, CMMD, PVM, and MPI

William Saphir (NAS) presents a detailed comparison of several widely used libraries. He focuses on performance, functionality, and design.

Parallel I/O on Highly Parallel Systems

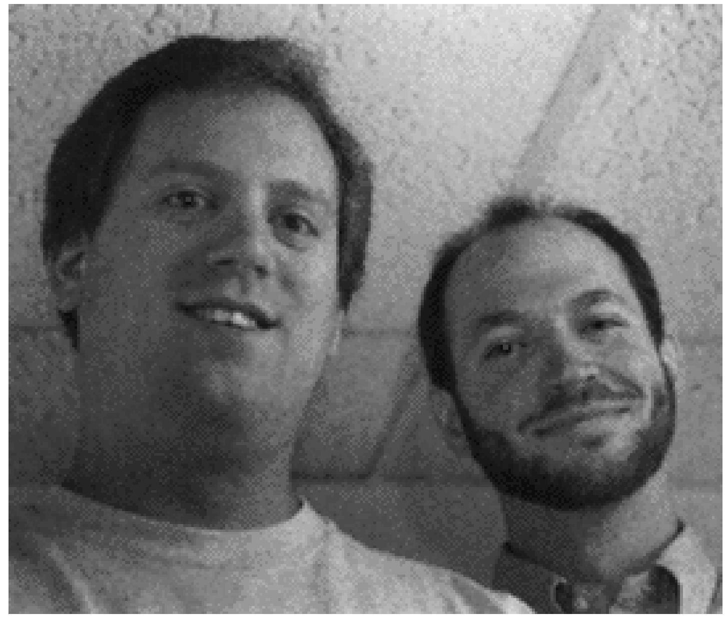

Samuel Fineberg (left) and Bill Nitzberg, both of NAS, compare the I/O performance of a traditional vector supercomputer with that of various highly parallel systems. The authors also present benchmarking techniques for evaluating current parallel I/O systems and techniques for improving parallel I/O performance.

Performance Tuning on RISC Systems

David Bailey (NAS), Charles Grassl (Cray Research Inc.); Ramesh Agarwal, Fred Gustavson, and Mohammad Zubair (all of IBM); and Dick Hessel (Silicon Graphics Inc.) focus on applications. The authors discuss the DEC "Alpha" (used in the CRAY T3D), the IBM Power2 (used in the IBM RS6000/590 and the SP2), and the MIPS/SGI R8000 (also known as the TFP).

Methodologies and Tools for Tuning Parallel Programs-Facts andFantasies

Jerry Yan (Recom Technologies/NASA Ames) reviews the fundamental concepts involved in analyzing and improving the performance of parallel and heterogeneous message-passing programs. Yan also contrasts some alternative strategies. His analysis is based on actual user experiences at NASA Ames Research Center. Yan discusses some limitations of these tools and methodologies, and outlines recent approaches taken by vendors and researchers to address these limitations.

Papers:

Here are brief descriptions of papers presented by members of the NAS Systems Division and some of their affiliates.

Non-numeric Search Results on the EM-4 Distributed-memory Multiprocessor

Andrew Sohn (New Jersey Institute of Technology and a JOVE fellow at NAS), Mitsuhisa Sato, Yuetsu Kodama, Yoshinori Yamaguchi (Electrotechnical Laboratory, Japan); and Shuichi Sakai (Real World Computing-Tsukuba Research Center, Japan) believe that "real" supercomputers should be able to effectively execute non-numeric problems. The authors present their experiences. See the article on Sohn's work, entitled "JOVE Program Researcher Works on Grid Partitioning for IBM SP2 at NAS."

Performance Evaluation of Three Distributed Computing Environments for Scientific Applications

Rod Fatoohi and Sisira Weeratunga (both of NAS) present performance results for the DCF and LACE clusters, and an Intel iPSC/860 machine, as well as performance results of several parallel algorithms for the three simulated applications.

NAS Experiences With a Prototype Cluster of Workstations

Karen Castagnera, Rod Fatoohi, Edward Hook, William Kramer, John Musch, Chuck Niggley, Bill Saphir, Doug Sheppard, Merritt Smith, Ian Stockdale, Shaun Welch, Rita Williams, and David Yip (all of NAS), Doreen Cheng and Craig Manning (formerly of NAS), present the one-year effort at NAS to implement a large, loose cluster of workstations, based on an existing Silicon Graphics Inc. pool of systems.

Applications Performance Under OSF/1AD and SUNMOS on the Intel Paragon XP/S-15

Subhash Saini and Horst Simon (both of NAS) present test results of performance and scalability.

An Efficient Abstract Interface for Multidimensional Array I/O

Marianne Winslett (University of Illinois and former guest speaker at a NAS New Technology Seminar) and Kent Seamons (University of Illinois) show how to produce advanced I/O libraries supporting more efficient layout alternatives for multidimensional arrays on disk and in main memory. The authors also present their experience to date in applying these techniques to CFD applications in the areas of checkpoint/restart, output data, and visualization.

A Portable Debugger for Parallel and Distributed Programs

Robert Hood (NAS) and Doreen Cheng (formerly of NAS) describe the design and implementation of a portable debugger. The debugger's design incorporates a client-server model in order to isolate nonportable debugger code from the user interface.

Improved Load Distribution in Parallel Sparse Cholesky Factorization

Robert Schreiber (NAS) and Edward Rothberg (Intel Corp. and guest speaker at a NAS New Technology Seminar) discuss a performance result of nearly 3.2 billion floating point operations per second that has been achieved with a technique that improves load imbalance.

Next Article Contents Main Menu NAS Home

CFD and NAS Information Demos Show `Big Picture'

by Elisabeth Wechsler

Two demonstrations at the NAS booth at Supercomputing '94 present the results and underlying concepts of work being done at NAS. The idea, according to Kevin McCabe, manager of the NAS Visualization Lab, is to educate conference visitors about computational simulation and other NAS research.

Overview of Computational Process

One demo lets visitors interact with software that gives an overview of the computational fluid dynamics (CFD) simulation process, including:

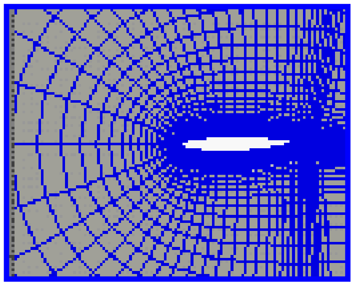
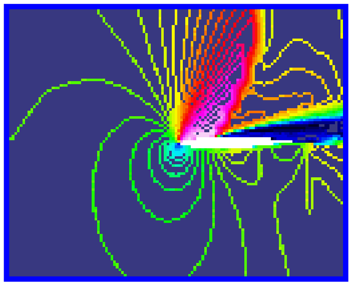
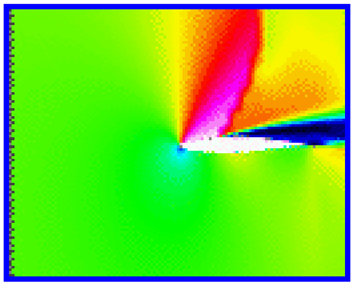

- construction of a two-dimensional airfoil and surrounding grid
- manipulation of mach number and angle of attack parameters to a flow solver
- selection of various visualization techniques (including contour lines, color-mapped surfaces, isosurfaces, and particle traces)

The intent of this demo is to give visitors a hands-on understanding of "a little of the discipline of computational simulation and to show why a supercomputer is necessary," McCabe said. He estimates that the average viewer will spend about five minutes interacting with the simulation.

NAS `Weblet' at a Glance

The "Information Assistant" demo features a condensed version of NAS information available on the World Wide Web, allowing users to browse at their leisure for various NAS offerings, using the popular

Mosaic viewer. Included on the "weblet" are excerpts of the 1992-93 *NAS Technical Summaries*, the most recent NAS Parallel Benchmarks results, the NAS Mission Statement in multimedia format, and animations of compressed image files in moving picture expert group (MPEG) format. Printed versions of the publications are also available from the NAS Documentation Center (send email to **doc-center@nas.nasa.gov**).

IofNewt Demo Shows Remote Wind Tunnel Experiments

The IofNewt (Integration of Numerical and Experimental Wind Tunnels) allows scientists to remotely observe and control wind tunnel experiments. This demo shows actual work in progress, and highlights the potential for huge research and development cost savings as well as dramatically shorter design cycles. The IofNewt project and a sub-element, the Remote Access Wind Tunnel, are highlighted in one of the three plenary sessions at Supercomputing '94.

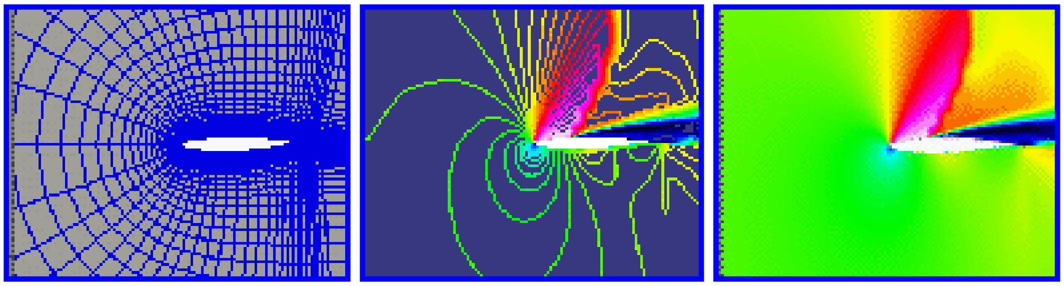

These illustrations show different parts of the CFD demonstration at the NAS booth at Supercomputing '94. The left figure shows points of intersection in an airfoil grid construction, where the flow solver calculates the fluid flow. The middle figure shows the fluid flow at a selected angle of attack and Mach number, illustrating how those parameters affect the aerodynamics of the airfoil. The right figure shows the selection of a cutting plane visualization technique to display the airspeed at all points around the airfoil.

to the article.

Stereoscopic Demo Presents Solutions to Aerospace Problems

by Elisabeth Wechsler

NAS captures the drama of stereoscopic viewing techniques for a demo at its Supercomputing '94 booth that presents "specific solutions to problems in aerospace," noted Kevin McCabe, manager of the NAS Visualization Lab. He explained that part of the NAS mission is to help U.S. industry solve aeronautics problems.

"Stereoscopic viewing can resolve visual ambiguities in the data, inherent in viewing three-dimensional (3D) information with a 2D screen," McCabe continued.

Solutions Inspired by NAS Research

The solutions shown in the demo were inspired by actual NAS research. They address such practical problems as maximizing aircraft fuel efficiency, shortening the design cycle, understanding acoustics, and improving geometric modelling techniques to enhance understanding of complex flows.

For the NAS demo, conference visitors wear "stereo" goggles in order to see different left and right eye views. To add impact, the demo data is being projected onto a 54-inch by 41-inch screen, with audio speakers playing background music in the separate, theater-like setting. The projector model used in the demo is an Electrohome Marquee 9000. The CrystalEyes goggles are manufactured by StereoGraphics Corp.

NAS demo teams for Supercomputing '94: from left, Chris Gong, Kevin McCabe, Dennis DeRyke (Stereoscopic demo); Sandy Johan and Steve Bryson (Virtual Wind Tunnel demo).

Demo Pinpoints Multiprocessor and Communication Code Errors

by Elisabeth Wechsler

A demonstration in the NAS booth of p2d2, a debugger currently under pre-alpha development, focuses on finding code problems peculiar to distributed computations. Specifically, classes of errors involving multiple processes across networks and deadlocked communications processes are featured, as well as techniques found in traditional sequential debuggers.

The idea of the demo, according to Robert Hood, of the NAS parallel tools group, is to give viewers the familiar format of a traditional debugger but one whose functionality has been extended to multiple processes.

p2d2 Tailored For CFD Researchers

"We're also paying attention to CFD code in the debugger architecture so that eventually we can offer CFD algorithm development support through the use of a visualization tool linked to the debugger," Hood said.

The first part of the demo shows code with a communication problem and how the debugger would be started and automatically attached to all processes participating in the job, he explained. The second part shows how the debugger would work if it were started with the program in the first place, in order to see what happens *before* a crash," Hood said. "The debugger needs to work in both of these modes in order to show how the tool can isolate the communication problem," he added.

Hood hastened to say that p2d2 is in a very early stage of development. With continued development he hopes to demonstrate how specific communication problems are solved.

In addition to Hood, Jeff Becker, also of the NAS parallel tools group, and Doreen Cheng (formerly of NAS) have worked on p2d2. Hood and Cheng are presenting a paper at Supercomputing '94.

Complements NTV Tool

If the debugger is started with the p2d2 (or, portable parallel/distributed debugger) program, the user can stop and look at errors as they are occurring, he said. The NAS Trace Visualizer (NTV) tool looks at the program after it has been run and produces traces that can be visualized.

"With p2d2, you can stop and look at ancillary problems as the program is running. With NTV, you can see the entire execution history but only for the trace data that was collected," Hood said, adding that he envisions CFD users employing both tools, since they are complementary.

- Descriptions of NAS tutorials and papers

Next Article | Contents | Main Menu | NAS Home

`Newsflashes' -- Updates on NAS Projects

Here are brief updates on several NAS projects previously reported in *NAS News*. Look for details on major upgrades to these and other projects in future issues.

The next release of **FAST**, the Flow Analysis Software Toolkit is being rewritten "from the ground up" and is currently in the prototype stage, according to Tom Whittaker, FAST team lead. The primary goal of FAST2 is the visualization of large, time-dependent datasets both on local workstations and in communication with more powerful computers. Major additions include sophisticated animation capabilities, a streamlined method to add modules, and a flexible user interface.

UFAT (Unsteady Flow Analysis Toolkit) developer David Lane reports that he and David Kenwright are working to improve the software's particle-tracing speed. Lane and Kenwright are creating new visualization techniques for particle traces, including: color contours on grid surfaces, cutting planes, isosurfaces, and stream surfaces. Lane estimates that development will be completed by April 1995. They are also creating an interactive distributed version of UFAT, which will compute on a Convex and send the results to Superglue (a programming environment combining LISP and C to implement interactive 3D graphics applications, being developed at NAS by Jeff Hultquist, Eric Raible, and Kristina Miceli) on a Silicon Graphics Inc. workstation for visualization.

NAStore, the NAS mass storage system, doubled the disk storage capacity allotted to users, from 800 gigabytes to 1.6 terabytes. The two-stage process was completed last month, according to Dan King of the NAStore group.

NTV (NAS Trace Visualizer) was released to users last April. A version for the IBM SP2 native trace is being tested at NAS and may be available for users to test this month, according to Louis Lopez, NTV developer.

JOVE Program Researcher Works on Grid Partitioning for IBM SP2 at NAS

by Elisabeth Wechsler

Andrew Sohn, assistant professor of Computer and Information Science at the New Jersey Institute of Technology, Newark, is the first research fellow to work at NAS under the NASA/University Joint Venture in Research (JOVE) program.

JOVE seeks to broaden research in space science and technology from the few hundred U.S. colleges currently participating. The three-year fellowships, which provide seed money and mentoring at established NASA centers for junior faculty, are designed to promote new research in NASA-related fields.

Algorithms for IBM SP2 at NAS

Sohn's research project involves dynamic grid partitioning for computational fluid dynamics (CFD) applications running on the IBM SP2 at NAS. Working closely with Horst Simon and Eric Barszcz, both of NAS, Sohn is investigating algorithms for partitioning and load balancing for dynamically changing calculations on massively parallel machines.

He spent 11 weeks last summer at the NAS Facility creating new algorithms and completing the framework for his research project. With an account on the SP2, Sohn will work remotely from New Jersey on project refinements and experiments. He is presenting a paper about his research results at Supercomputing '94. (*See "NAS Offers Tutorials and Papers at Supercomputing '94"*.)

Sohn received his BS (electrical engineering), MS, and Ph.D. (both in computer engineering) degrees from the University of Southern California. His work experience includes IBM, The Aerospace Corp., and Electrotechnical Laboratory (the Japanese National Laboratory) under a National Science Foundation fellowship.

For the next three years, Sohn will continue his NAS-based research while teaching and developing the introduction of aerospace topics into the curriculum of his engineering-oriented institute. The research projects he is currently working on with his students include developing a parallel programming environment for the 80-processor EM-4 distributed-memory multiprocessor using medium-grain multithreading, and parallelization and implementation of hard sequential problems on massively parallel machines.

Another part of the JOVE program requires fellows and their institutions to reach out to pre-college teachers and students, as well as to the broader community, using the space program to stimulate the imagination and to elevate public science literacy. Toward this end, the JOVE orientation retreat held at Cocoa Beach, FL, in July, arranged for Sohn and some 100 other participants to witness the launch of the Columbia space shuttle at Kennedy Space Center.

Under the JOVE program, a university must commit a minimum of 25 percent full time equivalent time per faculty member to the project in order to permit adequate research progress. Student support for the project is shared 50-50 between NASA and the university up to a maximum of $10,000 of NASA support.

Ames Establishes Network Links

Network links, facilitated by the NASA Science Internet Office at Ames Research Center, are established to match the research needs of the faculty research associate. NASA also provides funds for equipment essential to the implementation of the JOVE research project on the host institution's campus.

Support for Sohn's summer salary and travel is covered by NASA. Following an evaluation of the first summer research internships, the institutions are invited to submit three-year proposals for full program participation. The second and third summers provide an opportunity for Sohn and other research associates to focus on their projects full-time. A comprehensive annual report of each institution's activities is required at the end of each academic year.

Recruitment for JOVE Program

To recruit participants, NASA posts an announcement of the JOVE program each fall in the Commerce Business Daily and in the Chronicle of Higher Education, and hosts an orientation conference. By January 31, a university wishing to participate submits applications from candidates and a letter of commitment outlining that institution's faculty research time, student scholarships, matching summer salaries, and so forth. Institutions are notified of their acceptance by March 1.

For more information about Andrew Sohn's project, send email to **sohn@nas.nasa.gov**. For information about the JOVE program, contact N. Frank Six, JOVE Program Director, University Affairs Office, NASA/Marshall Space Flight Center, Mail Code DS01, MSFC, AL 35812; phone (205) 544 0997.

Andrew Sohn, the first research fellow under the JOVE program at NAS, is presenting a paper on his experience with the EM-4 distributed-memory multiprocessor at Supercomputing '94.

Virtual Reality Model Trains Eye Surgeons Without Risk to Patients

by Elisabeth Wechsler

A virtual reality surgical simulation model and a teleoperated micro-surgical robot for eye surgery are being developed collaboratively by the University of Auckland, New Zealand, and the biorobotics laboratory at Massachusetts Institute of Technology (MIT). When completed, the system is expected to radically change the way eye surgery is performed.

Preliminary research results were presented on July 21 at the NAS Facility as part of the New Technology Series, in a lecture, entitled "A Virtual Environment and Model of the Eye for Surgical Simulation," by Mark Sagar, of the University of Auckland.

Model Uses Isosurface Algorithm

The model incorporates an isosurface algorithm--a technique for constructing three-dimensional (3D) surfaces--that was developed by David Kenwright, a research scientist at NAS. Kenwright completed the algorithm while working on his doctorate at the University of Auckland, where he met Sagar. (Currently, Kenwright is developing particle-tracing algorithms for unsteady flow visualization.)

The Micro-surgical Robotic System

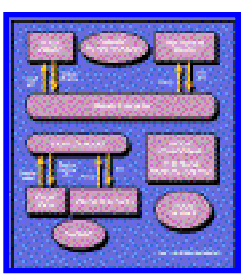

Sagar's virtual eye model is designed to work with a master-slave robotic system developed by Ian Hunter at MIT that can be controlled and monitored remotely: The surgeon works in a virtual environment and manipulates virtual surgical instruments attached to the master, or controlling, robot. These movements are transmitted via the master and slave computers to the slave robot, which holds the actual surgical instruments and reproduces the same movements, but scaled down as much as 100 times (*see Figure 1*). The scaling feature allows extremely precise procedures to be performed, Sagar explained.

Because micro-surgery is currently performed using hand-crafted tools, surgical intervention is limited by the dexterity of the surgeon and by the available workspace, he said. "One leading U.S. eye surgeon spends 30 percent of his time correcting the damage caused by other retinal micro-surgeons. There is a clear need in many areas of micro-surgery for more precise instruments and procedures."

Pressure Sensations Felt by Surgeon

Sagar explained that the pressure sensations experienced by the slave are reflected back to the master and can be felt by the surgeon. "The surgeon is holding an instrument that is mechanically attached to the master robot. It is through this connection that the surgeon feels the sensations," he said.

Forces measured by transducers in the slave are reflected back to actuators in the master. This feedback may also be scaled, allowing extremely delicate corneal tissue to be felt by the surgeon--perhaps for the first time-- offering the surgeon a new perspective of the procedure.

The surgeon receives stereo video images from cameras mounted on the slave. These are seen through either a head-mounted display or a rear projection screen. "The combined effect of the visual and mechanical feedback allows surgeons in the virtual environment to effectively see and feel an eye in front of them, whereas the actual patient may be in another room--or even in another country," Sagar said.

Computer-generated Environment

For training purposes and performance evaluation of surgeons using the system, a virtual reality surgical simulator forms part of the system. This way, "surgeons can get used to operating the system at no risk to the patients," he said.

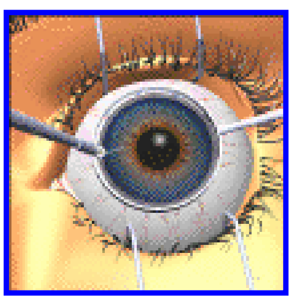

In simulation mode, the virtual reality system replaces the robotic slave subsystem: For example, the eye model and the virtual environment fit in the system as a module, which takes the place of the slave subsystem, allowing the same master equipment to be used in a new mode as a surgical simulator, he said. The stereo camera image produced by the slave is replaced by realistic computer graphics from a model of the eye (see Figure 2).

A mathematical technique called finite element analysis is used to calculate forces experienced by the virtual robot using a detailed model of the mechanics of the eye tissue. This way, both visual and mechanical simulation can be achieved, Sagar said.

The virtual reality environment provides the surgeon with additional information in a style similar to that used in a head-up display in which information is projected onto a transparent screen so that the surgeon can view it while simultaneously performing the procedure. The display may show planned cutting paths for radial keratotomy, with 3D arrows helping to guide the scalpel for exact placement.

Couples Visual, Mechanical Feedback

The physical system includes visual and tactile feedback. "The surgeon feels through his hands holding the virtual scalpel, which gives force feedback. The surgeon can wear a headset or look at a large 3D display, depending on which is more suitable for the operation," Sagar said.

The virtual environment includes a software framework that couples visual and mechanical simulations which can be run in parallel. The framework provides a relatively detailed level of realism for a virtual environment, and is a good testbed for evaluating new procedures and investigating ways of transferring more information to the surgeon during an operation, Sagar said.

Expensive Hardware Needed

"For the entire system to be operational--given the combination of graphics, finite element analysis, and robotics--the only limitation of the current system is the speed of the finite element analysis," Sagar said, adding that "approximately 5-30 gigaflops [billion floating-point operations per second] are needed to get the required analysis speed, compared with the 0.2 gigaflops available on our workstations."

Sagar explained that this technology lets surgeons monitor their own performance as they proceed--for example, hand tremor can be displayed in one corner, with the patient's vital signs displayed below. By recording these types of data during operations, surgeons can evaluate their performance and confer with colleagues about new methods through the virtual reality experience, he continued.

For more information about the virtual eye model, send email to **marks@ccu1.auckland.ac.nz**. For more information about other NAS training events, send email to **redmond@nas.nasa.gov**.

Next Article | Contents | Main Menu | NAS Home

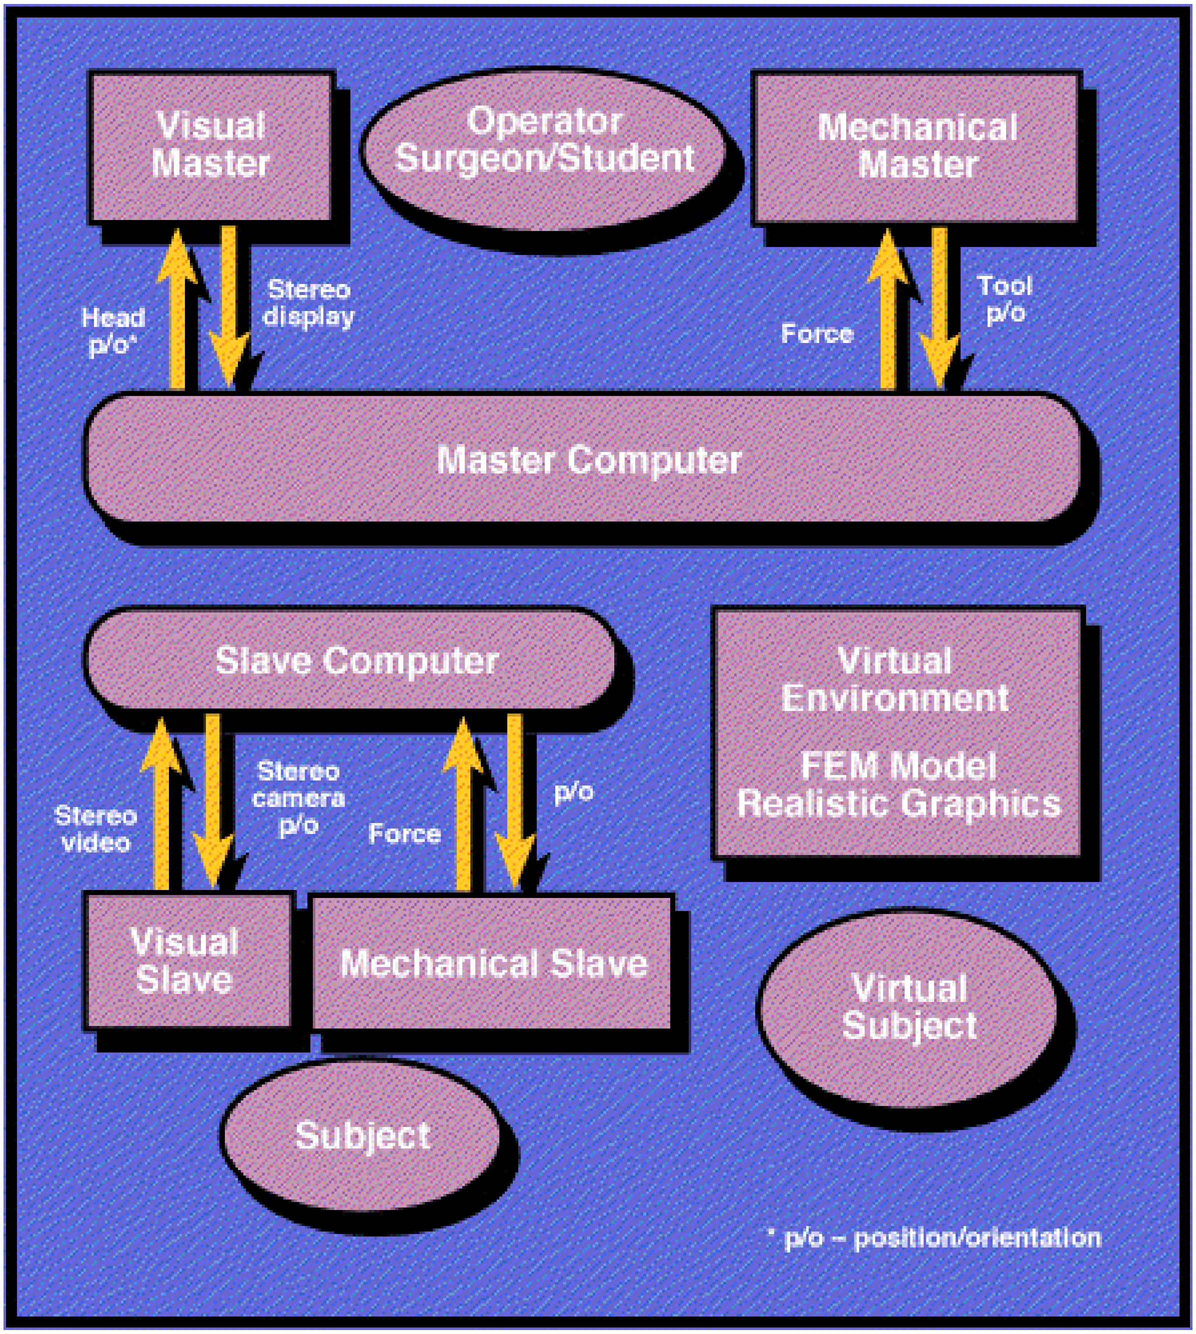

Figure 1. Block diagram of the micro-surgical robotic system, composed of mechanical and visual master and slave subsystems. The modular nature of the master-slave arrangement allows the images and forces fed back to the master from the slave to be repaced (or aided) by visual and mechanical computer simulations.

to the Main Text

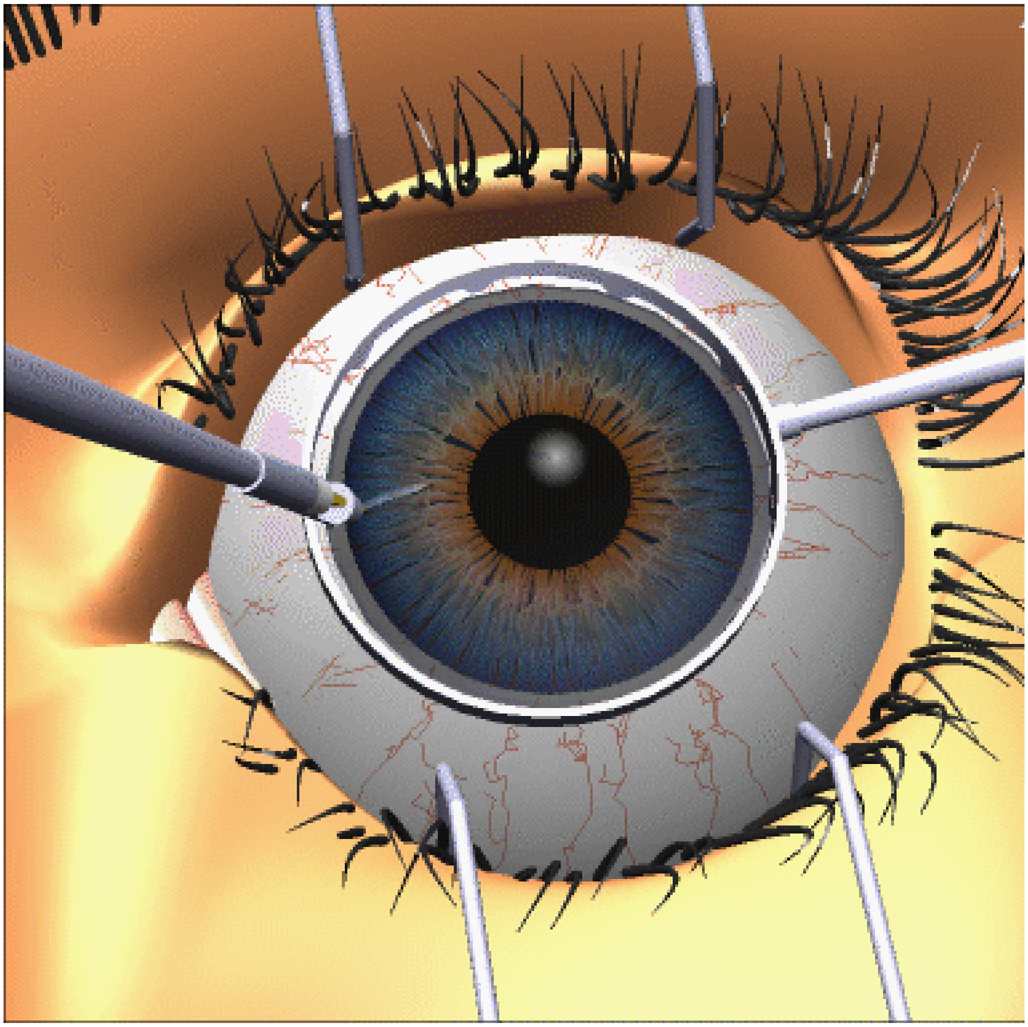

Figure 2. A scene from a virtual environment for eye surgery simulation shows the cornea being incised during radial keratotomy, in which incisions are made around the cornea to flatten it and thereby change its focusing properties. The resulting corneal deformations and forces being experienced are calculated from a finite element model of the corneal tissue. The simulator forms part of a teleoperated micro-surgical robotics system for eye surgery, a collaborative development effort between the University of Auckland, New Zealand, and Massachusetts Institute of Technology.

to the Main Text

NAS News Special Issue

Supercomputing '94--Vol. 2., No. 7

[Next Article] [Table of Contents] [NAS News Home Page] [NAS Home Page]

NAS NEWS

Volume 2, Number 7 Special Issue, Supercomputing 1994

NAS Special Edition for Supercomputing '94

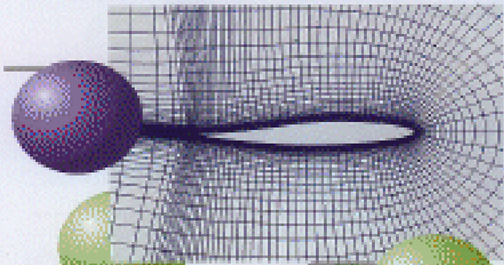

These illustrations show different parts of the CFD demonstration at the NAS booth at Supercomputing '94. Top figure shows points of intersection in an airfoil grid construction, where the flow solver calculates the fluid flow.

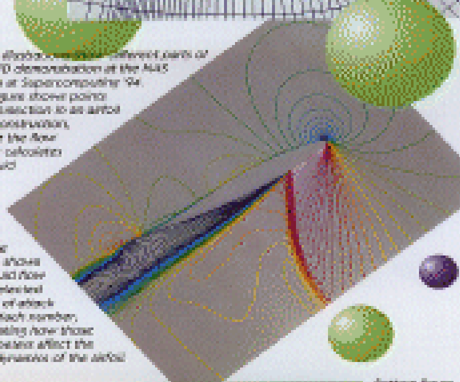

Middle figure shows the fluid flow at a selected angle of attack and Mach number, illustrating how those parameters affect the aerodynamics of the airfoil.

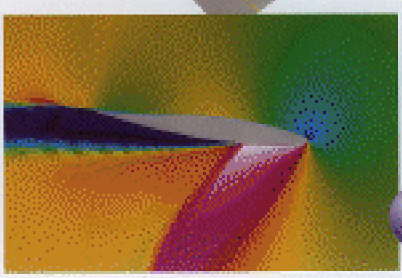

Bottom figure shows the selection of a cutting plane visualization technique to display the airspeed at all points around the airfoil.

CFD and NAS Information Demos Show 'Big Picture'

by Elisabeth Wechsler

Overview of Computational Process

NAS 'Weblet' at a Glance

More NAS Demos on Page 6